Teacher Answer Key

Case Study

1. California's BOT Disclosure Act requires that AI bots in commercial contexts disclose their artificial nature when sincerely asked. If Taylor had directly asked Alex whether they were human, and if the app falls under the law's scope, the app would have been legally required to disclose that Alex is an AI. Taylor never asked, partly because they did not know to — which points to a gap the law does not address: disclosure only when asked puts the burden on the user to know enough to ask the right question.

2. Taylor's situation falls on the deceptive high-stakes end of the spectrum, not the legitimate fictional roleplay end. Taylor was not engaging in a game with a known AI character — they believed they were forming a genuine friendship with a real person and shared deeply personal information on that basis. The reading identifies this kind of sincere, emotionally significant interaction as the context where disclosure matters most, because the harm from deception — loss of trust, feeling of violation — is greatest when the relationship felt real and important.

Multiple Choice

1. A
2. C
3. B

Matching

1-D, 2-C, 3-E, 4-B, 5-A

Fill-in-the-Blank

1) citations, 2) character, 3) accounts

Short Answer (Sample Responses)

Q1: The ethical weight increases when the person interacting with the AI is vulnerable and the interaction touches high-stakes decisions or emotional wellbeing. The reading gives two examples: the character.ai case in which a 14-year-old formed an intense emotional relationship with an AI persona and died by suicide, and mental health chatbots engaging with users in emotional crisis in ways that suggested human counseling. When an AI's deception affects a person's mental health, safety, or legal situation, the harm from that deception is qualitatively more serious.

Q2: The business rationale is that human-seeming interactions produce better customer satisfaction scores and reduce the friction users feel when they realize they are talking to a bot — making people more likely to complete their customer service interaction positively. The ethical concern is that this prioritizes commercial outcomes over honesty, using design to conceal the AI's nature in order to manipulate users' emotional responses, without giving them the accurate information they would need to make informed choices about how much to trust or disclose to the system.

Q3: Accept well-reasoned answers from either side. For proactive disclosure: the reading notes that the current system places the burden on users to know enough to ask, which Taylor's case shows is inadequate. People form genuine emotional connections based on a false premise, and requiring proactive disclosure would prevent that harm. For disclosure-when-asked: the reading notes that many legitimate uses of AI — roleplay, fiction, pedagogy — involve AI personas that are valuable precisely because they seem real. Requiring proactive disclosure in all contexts would undermine those uses. A well-argued response should engage with the specific evidence the reading provides.